

Beyond Predictive Algorithms in Child Welfare

Erina Seh-Young Moon
University of Toronto
Toronto, Canada
erina.moon@mail.utoronto.ca

Tegan Maharaj
University of Toronto
Toronto, Canada
tegan.maharaj@utoronto.ca

Devansh Saxena
Carnegie Mellon University
Pittsburgh, United States
devanshsaxena@cmu.edu

Shion Guha
University of Toronto
Toronto, Canada
shion.guha@utoronto.ca

ABSTRACT

Caseworkers in the child welfare (CW) sector use predictive decision-making algorithms built on risk assessment (RA) data to guide and support CW decisions. Researchers have highlighted that RAs can contain biased signals which flatten CW case complexities and that the algorithms may benefit from incorporating contextually rich case narratives, i.e. - the **casenotes** written by caseworkers. To investigate this hypothesized improvement, we quantitatively deconstructed two commonly used RAs from a United States CW agency. We trained classifier models to compare the predictive validity of RAs with and without casenote narratives and applied computational text analysis on casenotes to highlight topics uncovered in the casenotes. Our study finds that common risk metrics used to assess families and build CWS predictive risk models (PRMs) are unable to predict discharge outcomes for children who are not reunified with their birth parent(s). We also find that although casenotes cannot predict discharge outcomes, they contain contextual case signals. Given the lack of predictive validity of RA scores and casenotes, we propose moving beyond quantitative risk assessments for public sector algorithms and towards using contextual sources of information such as narratives to study public sociotechnical systems.

CCS CONCEPTS

• **Human-centered computing** → **Empirical studies in HCI**; • **Applied computing** → **Computing in government**.

KEYWORDS

algorithmic decision-making, algorithmic bias, risk assessments, child welfare, public sector

ACM Reference Format:

Erina Seh-Young Moon, Devansh Saxena, Tegan Maharaj, and Shion Guha. 2024. Beyond Predictive Algorithms in Child Welfare. In *Graphics Interface (GI '24)*, June 3–6, 2024, Halifax, NS, Canada. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3670947.3670976>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GI '24, June 3–6, 2024, Halifax, NS, Canada

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1828-1/24/06

<https://doi.org/10.1145/3670947.3670976>

1 INTRODUCTION

Predictive risk models (PRMs) have increasingly infiltrated our everyday lives, including within the criminal justice system [4, 32], homelessness [74], employment services [71, 112], the child welfare system [92, 109], and education sector [72], to study and predict the future given the presently available information and generate decisions that carry significant ramifications for individuals. While initially developed with the expectation that such systems can reduce costs and generate objective decision-making [19], recent research has become increasingly critical of their objectivity and impact.

The US child welfare (CW) sector is one such sector that has extensively adopted PRMs to predict child maltreatment risk. Given the high stakes for the lives of children and families, such algorithmic systems have been introduced to alleviate caseworkers' high workloads, minimize costs, and arrive at objective, evidence-based decisions with a particular focus on empirical risk measurement and prediction [15, 91, 92, 97, 100], i.e., proactively recognizing high-risk cases. However, recent research that examines stakeholder perceptions towards these systems has found that they use potentially biased quantitative inputs related to factors like poverty and unjustly surveil and stigmatize low-income and minority communities when they enter the CW system [44, 50, 86, 104]. There is also increasing awareness in the general scientific community on the relevance of negative results [47, 94, 114] to evaluate and audit the impact of algorithmic systems through multiple methodological lenses.

In recent years, scholars have increasingly explored how CW narrative data can be used to infer contextual insights around the CW system and be incorporated into PRMs. Saxena et al. applied topic modeling on caseworker casenotes to illuminate invisible patterns of frontline caseworker labor [95] and uncovered that while risk in CW is primarily measured as a function of *risk introduced by clients*, this does not fully account for a clients' holistic circumstances and the dynamic multiplicity/temporality of risk factors underpinning CW cases [94]. Field et al. [46] also found training computational models using text data **with** quantitative, structural administrative data (e.g., demographic information, past public welfare receipt history, criminal history) may mitigate racial disparities in out-of-home placements predictions. With these works in mind, we critically examined the predictive validity of the *conceptualization of risk* in predictive risk models prevalent in the child welfare sector (CWS). We quantitatively compared signals from RA data and CW

Model (best-performing in overall accuracy)	True Negative Rate
Random Forest using RA scores (AAPI scores)	0.167
Random Forest using Topic Model Probabilities	0.357
Constant Predictor	0.5
Random Predictor	0.2

Table 1: Classification model metrics for the top performing models: We show how often the top performing models could identify when children were not reunified with their bio-parent(s) given a balanced dataset. The Constant Predictor predicts reunified ('R') and not reunified ('NR') outcomes 50% of the time each. The random predictor predicts 'R' outcomes 80% of the time and 'NR' 20% of the time.

narratives to ask **RQ1**: How does risk measured in CW risk assessments affect case outcomes? **RQ2**: How do case narratives inform child welfare caseworker decisions and case outcomes? and **RQ3**: How can researchers better support child welfare caseworkers?

We partnered with a child welfare agency in a US Midwestern state and examined casenotes written by caseworkers and CW risk assessment (RA) scores assessed for 277 families discharged from the agency between November 2018, to March 2022. We trained 14 classifier models using support vector machines (SVM) and random forest models, incorporating narrative data from CW caseworker casenotes using topic modeling and RA scores from the Adult Adolescent Parenting Inventor-2.5 (AAPI) and North Carolina Family Assessment Score (NCFAS) (two commonly used RAs which aims to comprehensively measure case risk factors for families [35, 67]). We compared the predictive validity of our trained classifier models to examine how well casenote data and RA scores can predict reunification (discharge) outcomes. We also qualitatively examined topics in our topic model solution to better understand the types of case signals our topic model solution provided to the trained classification models. In summary, our study makes the following contributions.

- We quantitatively deconstruct AAPI and NCFAS scores, which measure child maltreatment risk across 100 metrics, to examine their ability to correctly predict discharge outcomes. We find quantitative support that RAs are largely irrelevant in informing discharge outcomes. We find that while the best performing classifier built using *only* RA scores shows a high accuracy rate of 86.4%, the model can correctly identify cases that end in non-reunification ('NR') between child and bio-parent(s) only 16.7% of the time (seen in Table 1). This is worse than a constant predictor or random predictor.
- Through computational text analysis, we show that case narratives provide different signals around cases compared to the RA scores. Case narratives can highlight contextual details between individuals involved in a case that can affect the trajectory and outcome of cases. We find that a random forest model built using *only* topic model solution data can correctly identify cases that end in 'NR' 35.7% of the time (see Table 1). This is better than a random predictor but worse than a constant predictor.

- We show how casenotes inherently contain limitations in prediction tasks in CW as they present the writer's perspectives. We also show narratives in aggregation can provide the contextual rationale behind CW decisions.
- We computationally critique CW algorithms that use RA scores to predict and affect case outcomes. We find support to shift away from predictive algorithms in CW and towards using contextual information such as narratives to study public sociotechnical systems.

2 ETHICAL CONSIDERATIONS

We had Institutional Review Board (IRB) approval from our university to use CW casenotes and discharge records for this research. The data we used for this research was bound by the IRB and a research agreement with the CW agency, which only permits us to use the data for research purposes. We are also aware that beyond the IRB guidelines and data-sharing research agreement, the nature of our research is highly sensitive, and the process in which CW data is collected is a morally complicated matter related to privacy, oppression, and surveillance. Considering the pervasiveness of predictive risk models (PRMs) in CW and their biased impact on families, we hope our study can illuminate the limitations of using RA scores to build and guide CW decisions. In this study, to minimize the risk of unintended harm, we considered the principles of beneficence, respect for law, and public interest when determining how and to what extent we should utilize the information from our datasets for our analysis. We anonymized names and numbers from the data to ensure private information does not leak into our predictive models and do not make any raw data public.

3 RELATED WORK

3.1 Algorithmic "Risk" Assessment Tools in the Public Sector

Public sector agencies in several countries are experiencing increasing pressures from policymakers to use cross-sector public administrative data to improve decision-making in managing and delivering public services by employing algorithmic systems [2, 56, 91, 110]. Predictive risk assessment algorithms are now being used to make high-stakes decisions about human lives by several public agencies such as the child welfare system, criminal justice system, public education, housing authority, and public benefits programs to decide which neighborhoods require more surveillance, predict the risk of recidivism [111], tackle school segregation and assign students to school zones [87, 88], determine student performance [113], and to assess the risk of long-term unemployment for a person [2, 56].

As Redden et al. [83] note, one of the critical drivers in this push for algorithmic systems for governance has been the preemptive recognition and mitigation of "risk", i.e. – identifying clients in the riskiest circumstances and extending services to them. However, many of these algorithms have inadvertently exacerbated racial biases where minority communities are being over-surveilled [111], schools are becoming more segregated [106], and minority families are being over-investigated for allegations of child maltreatment [54]. On the other hand, due to an overemphasis on clients "most in need", resources from public benefits programs are being directed away from citizens who only require temporary assistance to regain

stability in their lives, i.e. – clients who require significantly fewer government resources but are labeled “low risk” and are consequently, screened out of consideration [44]. As highlighted by prior work, four critical factors make the task of risk prediction especially hard in the public sector – 1) outcomes in the public sector such as the risk of recidivism and risk of child maltreatment are poorly and inconsistently defined [9, 44, 53, 92], and 2) it is not possible to directly observe and assess an outcome such as “risk of future child maltreatment”, therefore, algorithms are trained to predict proxy outcomes such as the likelihood of re-referral or placement in foster care within the next two years [25, 60]. Furthermore, 3) a person’s life can stabilize and destabilize several times which makes it hard to assess what may constitute a successful outcome or intervention [56, 91], and 4) administrative data in the public sector is poorly and inconsistently collected which leads to several algorithmic biases in outcomes that seek to predict client’s behavior and may lead to unfairness in decision-making processes [34, 51, 94]. Therefore, there is a need to critically examine how predictive risk models are formulated in the public sector and whether they are achieving better outcomes for citizens.

3.2 Predictive Risk Models (PRMs) in Child Welfare

Predictive risk models (PRMs) are deeply embedded in various stages of the CW system [91]. Initially introduced to mitigate concerns around subjective decision-making by individuals and to arrive at data-driven, evidence-based, consistent decision-making for CW cases [92, 97, 100], PRMs help determine which families should be screened in for further investigation when an allegation of child maltreatment is made, deciding whether to remove a child from their biological parents and predicting future maltreatment risk [44, 51, 91]. For a more in-depth analysis of child welfare algorithms, we direct the readers to Saxena et al. [92] which provides a comprehensive review of predictive risk models in CWS.

Recent scholarship studying CW algorithms has largely used a participatory approach to evaluate the downstream effects of implementing these systems on different racial or income groups, envision solutions to mitigate algorithmic harms [25, 52, 104], and qualitatively understand stakeholder perceptions and usages of these computational systems [57, 61, 92]. Notably, Stapleton et al. [104], and Brown et al. [18] highlighted the harms caused by these models finding that CW stakeholders (e.g., caseworkers, families involved with CWS) have concerns about how the models focus on case deficits (i.e., risk) rather than considering case strengths which can inadvertently perpetuate stigmatization and racial biases. Moreover, researchers have found that subjectivity and caseworker discretion are absorbed into PRMs despite their touted objectivity. Inputs into these models often involve factors that require inherent interpretation such as parent cooperativeness, parenting attitudes, or home visits [15, 91] and systematic factors where due to high staff turnover, inexperienced caseworkers measure risk based on impressions [15, 36, 96]. In field studies, Cheng et al. [25], Kawakami et al. [61], and Saxena et al. [91] found that caseworkers actively exercise discretion to mitigate racially biased outcomes from PRMs, gamed the models, or tried to guess how features affected risk scores to better serve their clients. Saxena et al. also [94] further

problematized empirical CW risk prediction. The authors found that PRMs often treat risk as a static construct when in fact, risk in CW cases is dynamic and caused by multiple factors, including the child welfare system itself which faces a shortage of good foster homes and experienced caseworkers who can better support clients. In sum, several qualitative [61, 91], quantitative [25, 96], and mixed-methods studies [94] conducted on PRMs in child welfare have questioned the validity of empirical risk prediction itself.

3.3 Assessing the validity of Child Welfare PRMs

Coston et al. [38] and Raji et al. [82] argue it is important to examine the validity of decision-making algorithms to evaluate their functionality. In a related vein, Saxena et al. [94, 96] recently raised construct validity concerns around PRMs in child welfare, finding marked divergences between how risk is quantified as a static construct in risk ratings and how they are depicted as a temporally dynamic construct in CW casenotes. These studies raised the need to better understand and account for contextual systematic and procedural factors to measure ‘risk’ in CW. And yet, interestingly, although PRMs mostly focus on predicting child maltreatment ‘risk’, to the best of our knowledge, few studies quantitatively examine their validity to question how ‘risk’ in these models is conceptualized, measured, and utilized to guide decision-making. To date, Vaithianathan et al. [108] have examined the predictive validity of a PRM by measuring the association between risk scores and future child hospitalization/emergency department visits in the follow-up period. However, because their study focused on probabilistic associations between risk scores and future maltreatment events, their findings could be subject to confounding factors as CWS would likely intervene once a family is deemed high risk.

In sum, we find that while many child welfare PRMs focus on measuring and predicting child maltreatment risk, there is a gap in the literature around our understanding of how risk is conceptualized in these models from a validity perspective. In this study, we specifically examined the predictive validity of risk assessments (RAs) used to build PRMs by examining how RAs and casenotes inform discharge outcomes. In CW, discharge outcomes signal whether the courts and caseworkers determine a child is safe to be reunified with their parents. Using machine learning techniques, our study asked how risk measured in RAs affects CW discharge outcomes and how casenotes inform CW decisions and outcomes. Through our research questions, we seek to better understand how researchers can better support CW caseworkers.

4 RESEARCH CONTEXT

We partnered with a CW agency in a Midwestern US state. The agency is contracted by the state’s Department of Children and Families (DCF) to support foster homes and offers case management services in compliance with all DCF directives. When allegations of child maltreatment are reported and hereafter substantiated by a DCF worker, the case is referred to the CW agency. Following DCF standards, caseworkers conduct initial assessments to determine the family’s needs and provide recommendations to the court. Then, after discussions with birth parent representatives, the judge, and the district attorney’s office, CW teams at the agency work together

to offer services to families. Within the CW agency, various teams provide different services to families, including teams that focus on offering in-home services, preventing sex trafficking, and helping foster youth transition out of foster care.

For this study, we specifically obtained casenotes and a discharge dataset with RA scores from the Family Preservation Services (FPS) team at the CW agency. When a case is assigned to the FPS team, FPS caseworkers assist bio-parents achieve reunification with their children by providing crisis support, parenting classes, and helping improve family functioning [26]. Following FPS intervention, FPS caseworkers provide evidence to the DA's office and judge to recommend a discharge outcome for the family, such as reunification with their bio-parent, placement into foster care, kinship care, or adoption. Discharge outcomes on families meaningfully signal whether caseworkers and the courts deem a child safe to be reunited with their bio-parent(s) or if they should be placed into other placement forms such as foster care or adoption [29].

An integral part of FPS caseworker duties include recording casenotes following stringent documentation standards guided by social work theory [22]. These casenotes serve multiple critical purposes, including, providing evidence for the DA's office to recommend reunification, ensuring accountability to caseworkers and the agency, and helping guide the child welfare process by providing a summary of services offered to families [22, 102]. DCF also requires FPS caseworkers to complete quantitative RAs to assess family risk and safety per the Child Abuse Prevention and Treatment Act [29]. In the CW agency's Midwestern state, caseworkers are required to conduct RAs at intake and discharge (when families enter and before leaving CW services) to consistently collect information on factors that threaten child safety. Based on the RA scores and casenotes recorded, caseworkers then use casenotes and RAs as evidence to justify case discharge recommendations.

5 METHODS

5.1 Dataset

For our study, we combined two different datasets provided by the Family Preservation Services (FPS) team at our partner child welfare agency, the discharge dataset and the casenote dataset.

Casenote dataset The casenotes dataset comprised 12,576 casenote records for 471 families that received services from the FPS team from May 1, 2019, to March 31, 2022. Families were not necessarily involved with the agency for this entire period but engaged with the agency within this timeframe. The dataset included all details pertaining to any caseworker interactions and observations of persons involved in a case (e.g., biological parents, children, foster parents) and internal communications between caseworkers. Interactions could be in the form of phone calls, email messages, and in-person or virtual meetings. Each casenote record included a family ID and narrative casenotes.

Discharge dataset The discharge dataset comprised 371 discharge records for 339 families discharged from the FPS team from November 2, 2018, to March 24, 2022. Each discharge record included a Family ID, date of admission and discharge from the program, and RA scores for the AAPI [11], and NCFAS [64]. These two assessments are two commonly used RAs in CW that fulfill state DCF requirements and are similar to other CW assessments that

categorize risk into various levels¹. The AAPI asks parents to rate their parenting attitudes related to child maltreatment from 1 to 10 where scores can be binned into low (8-10), moderate (4-7), and high risk (1-3) [43]. In contrast, the NCFAS is completed by caseworkers and asks caseworkers to rate multiple domains of family functioning related to a family's environment, parental capabilities, family interactions, family safety, child well-being, and readiness for reunification [63, 64]. NCFAS scores range from -3 to +2 and can also include not applicable ('NA') or unknown ('UK') [77]. Negative scores indicate challenges, 0 signifies the baseline/adequate level, meaning no interventions are needed, and positive scores mean strengths [64]. By design, the AAPI and NCFAS scores are taken twice as a structured means for caseworkers to compare intake and discharge scores and assess how safety and risk factors in cases have been addressed following CW intervention. However, it is not always possible to collect all RA data. Our dataset showed that some intake and discharge RA scores were missing because families refused to take the assessments or caseworkers could not contact the bio-parents as they were hospitalized, incarcerated, missing, or had moved to a different state. Other times, the caseworkers determined the case was a low-needs case and an assessment was unnecessary or if the family had recently been re-referred to the agency and their scores were already on file.

5.2 Preprocessing the datasets

To track each family's entire narrative casenote history, we collated all casenote records in the *Casenote dataset* (N=12,576) by each family ID ordered chronologically (oldest to latest). We cleaned and anonymized the casenotes by removing stopwords, punctuation, numbers, and names. After preprocessing the casenotes, we noted that many of the casenotes were lengthy, with collated casenotes averaging 3,299 words, with 53,303 unique words across the casenotes, and there being 38,748 words in the longest casenote.

In the *Discharge dataset*, each record detailed the discharge outcome of a case which could be either adoption, guardianship/independent living, reunification, and concurrent permanency plan². As we wanted to understand how RAs can meaningfully inform whether it is deemed safe for a child to be reunified with their bio-parent, we grouped the discharge outcomes into two: reunified ('R') and not reunified ('NR')³. After dichotomizing the discharge outcomes, we noticed more than 80% of cases ended in 'R.'

Combined dataset Finally, we combined the two preprocessed datasets by using Family ID as the unique identifier so that we could have one set of casenotes and one discharge outcome for each family ID. Due to how data was recorded at the agency and because some families were continuing to receive CW services, some family IDs that had an associated casenote did not have a

¹AAPI (Adult Adolescent Parenting Inventory) has been administered in over 2 million administrations [45], and NCFAS has been used in over 1,000 US agencies and in 20 countries [73].

²This is when caseworkers are working towards reunifying children with birth parents while simultaneously identifying and working towards other placement options for the child such as guardianship with a relative or adoption in case placement with birth parents is unsuccessful [28]

³We categorized adoption, concurrent permanency plan, and guardianship/independent living outcomes as 'NR' and reunification outcomes as 'R'. We categorized concurrent permanency plan as 'NR' as this outcome signaled caseworkers anticipated a possibility that children would not be reunified with their biological parent(s).

Dataset Name	N	Features Included	Target Variable
TM	277	19-topic model solution probabilities (19TM)	Discharge outcome
AAPI	145	Intake & discharge AAPI scores	Discharge outcome
NCFAS	221	Intake & discharge NCFAS scores	Discharge outcome
AAPI + NCFAS	133	Intake & discharge AAPI and NCFAS scores	Discharge outcome
AAPI + TM	145	Intake & discharge AAPI scores and 19TM	Discharge outcome
NCFAS + TM	221	Intake & discharge NCFAS scores and 19TM	Discharge outcome
AAPI + NCFAS + TM	133	Intake & discharge AAPI and NCFAS scores and 19TM	Discharge outcome

Table 2: Descriptions of the six datasets (N here indicates the number of observations in the dataset)

discharge outcome and vice versa. In this case, we excluded those family IDs, so that we ended up with a collection of 277 discharge records with associated 277 collated casenotes for 277 families.

5.3 Topic modeling on narrative casenotes

We adapted the steps taken by Saxena et al. [95] and applied Mallet’s implementation of LDA topic modeling [14] on our collated and preprocessed casenotes for 277 families (from section 5.2). LDA can generate topic distributions over words for topics, produce distributions of topic probabilities for each text document, and scale to dense texts [5, 13, 76, 81]. Using LDA, we ran a 19-topic model solution, where we computed 19 topic probabilities for each of the 277 collated casenotes. Prior work by Baumer et al. [10] has argued that conducting grounded thematic methods can help illuminate thematic patterns in text from topic model solutions. Accordingly, to contextually understand topics generated from our topic model solution, the first author of this paper then used an open-coding process to manually inspect casenotes with the highest probabilities associated with each of the 19 topics and keywords associated with the topics to understand and label each topic [16]. After labeling each topic, we grouped topics together if they carried similar thematic content to better gauge the high-level themes that emerge from the casenotes (as seen in topic groupings by theme in Fig 1).

5.4 Dataset Creation

To compare and contrast how different combinations of RA (risk assessment) scores and narratives can inform discharge outcomes, we created six different datasets from the combined dataset outlined in section 5.2. Each dataset contained a combination of intake and discharge RA scores or topic model solution probabilities as features and the discharge outcome for the respective family ID. Table 2 shows the names and data included in the datasets.

5.5 Classification Models

We trained classification models using the six different datasets presented in Table 2 to understand how RA metrics can inform CW discharge decisions. We one-hot encoded all intake and discharge AAPI and NCFAS scores in the datasets and then created train and test sets for our datasets using a 7:3 ratio. Training and evaluating risk assessment classification models often carry challenges due to class imbalances common in risk assessment contexts. Class imbalances can cause models to fail to identify patterns from the minority class and instead favor the majority class [48, 59, 69].

Because our datasets were relatively small and imbalanced (more than 80% of cases in the datasets ended in ‘R’), we thus resampled and balanced the training datasets to include 1000 rows of training data where 500 observations ended in ‘R’ and 500 observations ended in ‘NR’ using synthetic minority oversampling techniques (SMOTE) [69]. SMOTE is useful because its variants can be applied to both continuous and categorical data [69], and for each minority class observation, the technique can quickly generate new examples by randomly finding closest neighbors in the feature space [24]. Next, we reduced the dimensionality of our datasets by applying principal component analysis (PCA) to denoise the datasets [79]. Finally, using the resampled and dimension-reduced datasets, we trained and compared the performance metrics of random forest models setting a threshold at 0.5 [17] and support vector machines (SVM) using a radial basis function kernel [37] for each of the datasets. We trained classifiers using random forests and SVMs because they can work with high-dimensional data efficiently and on problems that are non-linear [20, 68, 80]. While studies often adjust decision thresholds and validate risk assessment classifiers by inspecting the AUC of the classifier across different thresholds, Lobo et al. [70] discussed by Gerchick et al. [50], and Kwegyir-Aggrey [66] find (1) model comparisons with AUCs obfuscate different error types from different thresholds, (2) determining thresholds based on AUC ignores the inherently normative task of assigning costs to incorrect classifications, and (3) AUCs provide model performance summaries over the receiver operating characteristic (ROC) curve that may not be of relevance. In light of these findings, we evaluated the performance of the classifiers by inspecting the accuracy, false positive rate, false negative rate, and specificity rate of each of the models.

6 RESULTS

In this section, we present our findings from the trained classification models and organized by our research questions.

6.1 Risk assessments do not inform discharge outcomes (RQ1)

As seen in the orange rows in Table 3, we found that while classifiers built using only RA scores had high accuracy rates, the models were predominantly predicting children would be reunified with their bio-parent(s) and could not distinguish between cases where children were not reunified with their bio-parent(s). A false positive rate (FPR) in the predictive models indicates how often the models are predicting that a case would end in ‘R’ (reunification) when

Dataset	Model	Accuracy	FPR	FNR	Specificity
AAPI	SVM	0.864	1	0	0
AAPI	Random Forest	0.864	0.833	0.026	0.167
NCFAS	SVM	0.836	1	0	0
NCFAS	Random Forest	0.806	0.909	0.054	0.091
AAPI + NCFAS	SVM	0.825	1	0	0
AAPI + NCFAS	Random Forest	0.825	1	0	0
TM	Random Forest	0.702	0.643	0.229	0.357
TM	SVM	0.643	0.714	0.286	0.286
AAPI + TM	SVM	0.841	1	0.026	0
AAPI + TM	Random Forest	0.841	1	0.026	0
NCFAS + TM	SVM	0.806	0.909	0.054	0.091
NCFAS + TM	Random Forest	0.806	1	0.036	0
AAPI + NCFAS + TM	SVM	0.825	1	0	0
AAPI + NCFAS + TM	Random Forest	0.800	1	0.030	0

Table 3: Performance metrics for trained classifiers using for the seven datasets. The table shows the datasets used to train and test the model (1st column); the type of machine learning algorithm used (2nd column); and the accuracy, false positive, false negative, and specificity rate of the models (last 4 columns). Note that all datasets are class-balanced prior to training.

actually, the child was not reunified ('NR') with their bio-parent(s) in real life⁴. As seen in Table 3, SVM models trained using the AAPI or NCFAS or both AAPI+NCFAS dataset all had a FPR equalling 1. This meant that all cases in these models were miscategorized where all 'NR' cases (actual outcome) were predicted as 'R' cases. The random forest models trained using different combinations of the RA scores, similarly showed a high FPR, suggesting that RA data was not helping predict reunification outcomes regardless of the machine learning algorithm used. A false negative rate (FNR) in the predictive models indicates how often the models are predicting that children will not be reunified with their bio-parent(s) when in fact, they were reunified⁵. We noted that models built on RA data made almost no errors or any errors when predicting 'R' cases (FNR equaled 0 or nearly 0), meaning that the models were unable to distinguish cases that ended in 'R' and 'NR' through the RA scores. **Collectively, we noted classifier models that were trained using either AAPI (parent perspective) or NCFAS scores (caseworker perspectives), or both AAPI and NCFAS scores, did not improve the models' ability to correctly predict discharge outcomes. These results suggested RA scores were not identifying risk factors that determine case outcomes and were providing misleading signals regarding case risk factors.**

6.2 Case narratives can provide some signals towards discharge outcomes (RQ2)

We also trained classifier models using the probabilities from our trained topic model solution as features to examine whether narrative data can predict discharge outcomes. The blue rows in Table 3 show the model metrics that use only topic model data. A specificity rate in the models indicates how often the models are correctly predicting children are not reunified with their bio-parent(s) given the

fact they were actually not reunified⁶. Table 3 shows that classifier models that used topic model data had a lower accuracy rate, lower FPR, and higher FNR compared to models built solely on risk assessment data. Compared to models built using RA data, the lower FPR of 64.3% (random forest model) and 71.4% (SVM model) indicated that the models were miscategorizing proportionally fewer actual 'NR' cases as 'R', and the comparatively higher specificity rates showed that the models were doing a better job at predicting 'NR' cases. However, we also noted that models built using TM data had a FNR greater than 0, indicating the models were making several 'NR' predictions for actual 'R' cases. **These results suggested that topic model solution probabilities may be providing some indicators that can predict discharge outcomes (particularly in identifying and predicting cases that do not end in reunification) compared to using only risk assessment scores.** Our results resonate with Saxena et al. [96] who found differences between RA ratings and risks noted in caseworker casenotes, noting caseworkers were recording different facets of the same case in casenotes compared to RAs.

As shown in the gray rows in Table 3, we found that predictive models built using topic model data and RA scores exhibited similar performance metrics to models built using only RA scores. **Adding risk assessment data to topic model data did not improve the classifiers' ability to correctly predict 'NR' outcomes, further suggesting that while our topic model solution probabilities may be providing potentially more useful signals towards case outcomes, these signals are weak.** All in all, our results suggested topic model data may not be appropriate to predict case outcomes but could provide a unique lens into child welfare cases compared to risk assessment scores.

⁴FPR = FP/(FP+TN)

⁵FNR = FN/(FN+TP)

⁶Specificity is calculated by TN/(TN+FP)

6.3 Case narratives provide a holistic lens of inquiry towards discharge outcomes (RQ2)

To further understand the signals provided by the topic probabilities used in the classifier models, we manually examined original casenotes that had the highest probabilities associated with the topics in our topic model solution. **Manual inspection of casenotes showed that all casenotes were primarily written in a neutral tone that described facts and observations of case details (in fact, caseworkers in the US are trained to write casenotes in a neutral tone [22]). At the same time, we noted these casenotes revealed specific case details through the lens of caseworkers as they conduct a diverse range of street-level discretionary work.** These findings added support to findings from Section 6.2 that case narratives were providing contextual information on cases. As the main purpose of this study was to understand how RA scores and topic model solutions can inform case discharge outcomes, we present our topics by high-level themes to highlight how certain topics shared common themes in a manner consistent with prior literature that have performed topic modeling on CW casenotes [94, 95]. The following paragraphs further explain these themes. We also provide paraphrased exemplar sentences related to these themes.

Theme 1 was about case details that arise as caseworkers support birth parents to achieve reunification with their children by improving and strengthening relationships with their children, helping them understand their responsibilities, and fostering teamwork between different parties (i.e., foster parents, friends and family of bio-parents). Specifically, this theme revealed case details that emerged as caseworkers helped biological parents and foster parents establish boundaries and responsibilities (topic 5) [12], mediated interactions between children and birth parents when emotions ran high (topic 12), and observed and aided interactions between bio-parents and children and between siblings to improve family relations and placement stability (topic 6 and 9).

Topic 5 paraphrased sentence: [Bio-parent] has not been able to hold down a job for more than a month. [Caseworker] spoke to parent about how important it is to have a stable job and what can happen if she does not.

Topic 6 paraphrased sentence: [Baby] was making noises and bouncing on the jumperoo and [bio-parent] sat down in front of the baby. [Caseworker] encouraged [bio-parent] to talk to the baby and explained that babies love receiving attention and love from their caregivers.

Theme 2 highlighted information that arises when caseworkers communicate regularly with bio-parents and with other caseworkers within the agency to make collaborative decisions on cases. Case information emerged as caseworkers regularly attempted to contact bio-parents regarding case updates, schedule parenting classes and supervised visits, and ensure families get the support they need (topic 1) [12]. We found caseworkers also regularly internally shared details on cases with other caseworkers at the agency to make collective case decisions and coordinate services for families (topic 2).

Topic 1 paraphrased sentence: [Caseworker] tried to get in touch with [bio-parent] but the bio-parent did not answer the phone so caseworker left a voicemail.

Topic 2 paraphrased sentence from an email: Hello team. Thanks for providing the update. I am going to try to reach the [bio-parent] for their five-day visit and let the team know when we are scheduling the visit.

Theme 3 showed how caseworkers manage and arrange for catered services to be delivered to their clients. Topics under this theme detailed the paperwork caseworkers helped clients fill out and admin consent forms caseworkers needed to collect from bio-parents so that the children could receive required services such as medical services (topic 3) [105] and coordinate travel arrangements for children to supervised visits (topic 11). This theme also included details on how caseworkers addressed scheduling, or familial conflicts between biological parents or caregivers to ensure visits with children were taking place (topic 10) [103].

Topic 10 paraphrased sentence: [Caseworker] talked to [bio-parent] regarding [child] and visitation. [Bio-parent] said that bio-parent tried to get in touch with [foster parents] about her kids but was not able to get in touch with them.

Topic 11 paraphrased sentence: [Foster parents] called [caseworker] to determine if there was a visit today. [Caseworker] confirmed there is one today and had discussed this with [bio-parent] about the visit last week regarding visit details.

Theme 4 highlighted the different types of services CW staff were delivering to different families in addition to transportation and legal support. Casenotes revealed how caseworkers helped families obtain economic or material resources by connecting them to employment training centers, finding kitchen appliances and toys (topic 13), and checking in with biological parents or caregivers to ensure they were keeping up with a child's medical schedule (topic 17). Depending on family needs, this theme highlighted how caseworkers also drove children to supervised visits (topic 14), facilitated virtual interactions due to COVID-19 (topic 19), and advocated for families in in-person or virtual court sessions (topics 8 and 19).

Topic 13 paraphrased sentence: [Caseworker] will help [bio-parent] secure a weight scale. [Caseworker] left a voicemail to [furniture shop] about the delivery of bio-parent's stove and fridge.

Topic 17 paraphrased sentence: In today's visit, the caseworker noted no safety concerns. [Bio-parent] is up to date with all of [child]'s medical appointments.

Theme 5 was about caseworkers noting potential risks and safety concerns in cases and mediating them with protective factors. Caseworkers observed and recorded how children were dressed and acted before, after, and during transporting children to visits (topics 4 and 16), how biological parents responded to course content and questions during parenting classes (topic 7), and how children and birth parents interacted with each other (topic 15) [27]. Caseworkers also recorded the safety and suitability of a home for a child by observing who lived in the house and the presence of appliances, furniture, toys, and food at home (topic 18).

Topic 7 paraphrased sentence: [Caseworker] met with [bio-parent] online and completed a chapter of the Active Parenting Curriculum. [Bio-parent] was engaged in the meeting and it was clear bio-parent had read the material before the meeting.

Topic 15 paraphrased sentence: [Child] begins to rock in the chair and hitting the chair on the wall. [Bio-parent] tells child to stop and child

hits the wall harder. [Bio-parent] moves the chair away and the child begins rocking in the chair slowly.

7 DISCUSSION

Abebe et al. [1]’s “Roles for Computing in Social Change” outline four ways computational work can address social issues. Per the researchers, computing as a *diagnostic* can help us understand social problems that manifest in sociotechnical systems and as a *formalizer* to clarify how rule-based models define social issues via the inputs and outputs of the models. Computing can also offer a *rebuttal* to illuminate technical limitations and a *synecdoche* to highlight existing social problems by bringing renewed focus and attention to the issues. Our study takes on these roles and provides a quantitative commentary on existing literature on CW algorithms to affect change in how researchers study and support work in the child welfare domain and more broadly in the public sector.

7.1 A quantitative lens into the gaps in risk assessment data in predicting case risk factors (RQ1)

The poor predictive ability of our trained classifiers from section 6.1 acts as a *diagnostic* [1] of child welfare decision-making algorithms by raising questions about using RA data to build such models [46, 91, 92]. The high false positive rate (FPR) in Table 3 quantitatively showed that RA data (i.e., using only AAPI, only NCFAS, or AAPI + NCFAS scores) could not help predict discharge outcomes for cases. Most times, the models predicted ‘R’ (reunified) for most cases and nearly all cases that actually ended in ‘NR’ (not reunified) were incorrectly predicted to end in ‘R.’ The inability of the models to correctly identify ‘NR’ cases suggested that the RA scores were not signaling higher risk for these cases. These results were particularly notable because discharge AAPI and NCFAS scores were included in our datasets to train the classifier models, indicating that even discharge scores collected just before a family is discharged from the program cannot provide meaningful signals to predict case outcomes. These findings suggest that the pre-defined risk domains measured in the RAs flatten a case’s complexity [15] and cannot fully capture pertinent risk factors in cases, raising questions on whether we should be using RA data to inform child welfare outcomes.

The preponderance of the trained classifier models (built with RA data) in predicting ‘R’ outcomes for cases (shown by the FPR which equals or nearly equals 1 in section 6.1) find support for studies highlighting how different caseworkers can give very different risk ratings for the same case depending on an individual’s age, experience, stress level, and personal biases [15, 84, 91, 96]. It is possible that the motivations of the biological parents and caseworkers who completed the RAs may have influenced them to give scores that favorably align with their objectives. Most times, biological parents want to be reunified with their children, and the caseworkers are part of the Family Preservation Services (FPS) team, whose main objective is to support families achieve reunification as well [95]. Previous research on caseworker interactions with CW decision-making algorithms has shown that caseworkers will often adjust algorithmic inputs to serve their needs. For example, due to

limited good foster homes and to dissuade foster parents from ending placements to children, caseworkers often adjusted algorithm scores so foster parents can receive higher compensation for their services [91]. Furthermore, as AAPI and NCFAS were designed to measure theoretically derived risk-related domains from child welfare literature, caseworkers are likely working on helping families make changes to familial issues, which are constructs measured in the RAs [35, 64]. In this case, RA scores may be merely measuring the biological parent’s responses to caseworker interventions [96]. Jacobs and Wallach [58] argue harms related to fairness can arise in computational systems where there exists a mismatch in the intended construct to be measured and its operationalization. Consistent with this argument, our findings question the construct validity of risk prediction models such as the Allegheny Family Screening Tool (AFST) which seeks to predict child maltreatment using proxies such as child welfare records, demographic information, or recent incarceration records [30, 107], records which in themselves are influenced by potentially biased human discretionary decisions [44].

Prior research examining perceptions surrounding predictive risk models used in child welfare through interviews with affected stakeholders (e.g., family members, caseworkers, and specialists) [18, 91, 93, 104] have found that stakeholders expressed doubt in the utility of CW algorithmic decision-making systems as they primarily focus on the deficit of cases without accounting for system-level risks. In our study, we find quantitative support for these concerns by showing how RAs can present a reductive representation of complex cases and produce biased scores that are largely irrelevant to discharge outcomes while ignoring structural causes of social issues [62, 83]. A *diagnostic* [1] can serve as a starting point for social change by raising awareness of a problem. Through this study, we add quantitative support to existing qualitative concerns surrounding the use of RAs in algorithmic decision-making tools to affect child welfare case outcomes.

7.2 Shifting away from predictive child welfare predictive analytics by investigating case narratives (RQ2)

Our results from Section 6.2 show a higher specificity rate for classifiers built using only the TM dataset compared to models built using only RA data. As a *formalizer* [1], this suggests that compared to RA data, casenotes provide different signals surrounding cases and indicators that can correctly identify ‘NR’ outcomes. A closer dive into the topics in our topic model solution outlined in Section 6.3 highlights that unlike the focused risk domains measured by the RAs, the topics presented nuanced details on interactions between multiple agents that can influence the trajectory and outcome of cases. However, as a *rebuttal* [1], the lower than 50% specificity rate of the classifiers to correctly predict ‘NR’ outcomes in Section 6.2 raises caution against using topic models for predictive tasks. We noted that when we transformed topic model probabilities into continuous features to train classifier models, this step inevitably obfuscated much of the nuanced details recorded in narrative texts. Furthermore, manual inspections of casenotes highly associated with the topics outlined in Section 6.3 showed that although caseworkers often wrote detailed casenotes in a neutral tone based on

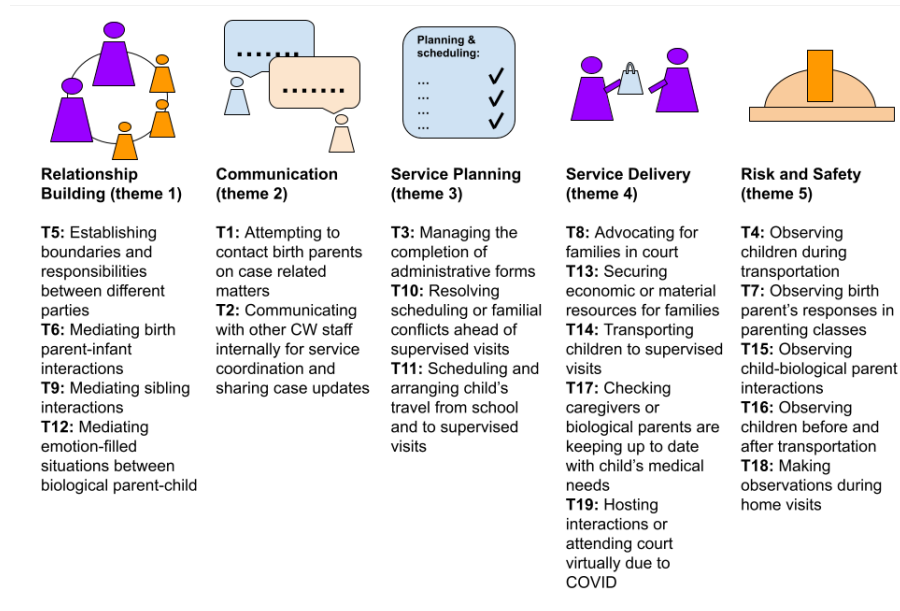


Figure 1: 19 topic model solution grouped by 5 high-level themes ("T" stands for topic).

observations and facts, they were still contextualized *caseworker perspectives* that may be biased because they could choose to include or omit specific case details. For example, we found different caseworkers described transporting children to bio-parent visits differently; one could record emotions exhibited by children during transportation, while others focused on the logistics.

Consistent with findings by Saxena et al. [94] and Field et al. [46], applying textual data for prediction tasks can therefore be problematic as that would mean relying on individual casenotes for prediction. Casenotes are not uniform nor the ground truth. They record socially situated experiences where different participants in the interactions might experience the same events differently due to structural power asymmetries [36, 40, 41, 85, 86]. Roberts [86] notes that police reports can entail contradictory records on the same CW case where one officer records a house is in fair condition with food while another records the same home is dirty with no food. In a similar vein, caseworkers may omit or underplay the oppression, surveillance, and coercion experienced by biological parents that fundamentally unpin CW interactions [94]. Furthermore, we noted that casenotes do not include all case details that may be critical to a case due to legal reasons. For example, we found that when a case involved an ongoing criminal investigation, caseworkers recorded that they omitted details following privacy laws [102].

Despite limitations on applying topic modeling on casenotes for predictive tasks, we identify key opportunities for using computational text analysis on casenotes. Nguyen et al. [78] argue that compared to smaller units of texts, texts in the aggregate can often unveil high-level themes. While caseworkers may exhibit different writing styles, our findings suggest that topic modeling on the casenote corpus can be a useful exploratory tool to uncover the dynamic and temporal dimensions in CW cases where family relations are complex and not self-isolated and exogenous organizational,

legislative, and policy tensions pose systemic risks [44, 97]. A common theme that emerges from HCI research that studies how CW workers use CW algorithms is how workers do not only rely on decision-making algorithms because they deem contextual knowledge of cases [60, 61, 91], and theoretical knowledge of trauma [91] with an awareness of organizational or placement constraints [90, 91] as critical factors that affect case decisions. Consistent with Saxena et al. [95], our findings show that child welfare narratives can unveil these contextual factors mentioned above and reveal invisible labor patterns and day-to-day power dynamics between stakeholders of child welfare. Understanding such details can be helpful for developers and researchers to examine how we can better support caseworkers in managing heavy caseloads and improve the quality of services offered to families [65].

7.3 Implications for algorithms in the public sector (RQ3)

Serving as a *synecdoche* [1], our research brings attention to the long-standing tenet in child welfare, which is: striving to protect children from maltreatment by assessing child maltreatment risk [21] without fully accounting for systemic risks [86]. Through a *diagnostic* [1], we quantitatively deconstructed RA scores and casenotes, finding that pre-defined risk parameters cannot correctly predict case outcomes, nor are casenotes suitable for predictive tasks in CW. From these findings, we add quantitative support to existing scholarship on decision-making algorithms used in CW [25, 89, 91, 92], the criminal justice system [31, 42, 58], and unemployment sector [3, 56, 98] which have been criticized for using proxies that violate construct validity [75], focusing on risk factors in cases, and leading to outcomes that are biased by disproportionately affecting historically marginalized population groups [4, 44].

Through the topics uncovered in our topic model solution outlined in section 6.3, we identify opportunities for using computational text analysis to investigate and support workers in public sector sociotechnical systems. The topics outlined in section 6.3 showed that aggregated narrative texts could unveil motivations, lived experiences, and work practices of stakeholders – holistic details not provided in RA data [95]. While individual narratives carry inherent limitations wherein they detail potentially biased *perspectives* [39], on the flip side, aggregated *perspectives* offer a unique, interpretive lens into contextual lived experiences that form a part of the complexities of sociotechnical systems, including invisible labor patterns, strengths in cases, and daily power asymmetries between stakeholders [95].

The nature of public sector work involves frontline workers making discretionary decisions which are often influenced by normative and socially constructed judgments regarding risk and safety [15, 55, 56, 96]. For example, in CW, caseworkers conduct home safety inspections to determine if a home is safe for children based on culturally influenced indicators of safety and neglect (e.g., cleanliness of a home, presence of furniture) [33, 49]. As the agency trains caseworkers to include observations of the home that *back up* their safety assessments in casenotes [22, 49], we argue narratives can provide a contextual rationale for decisions made. Similar to the child welfare sector, collecting documentation is a key task in the public sector, including job placement centers and the criminal justice system [3]. Here, records pertaining to cases record the dynamic negotiations between lawyers and district attorneys in criminal cases [31] and memos between job placement caseworkers and unemployed individuals [3]. Therefore, extending our study's opportunities for computational text analysis, we propose harnessing NLP techniques on narrative texts when investigating complex public sector sociotechnical systems. Our topics uncovered in section 6.3 suggest using narrative texts as a data source can help *humanize* actors in public sector domains and help us critically and holistically understand how computational or technical interventions may or may not support workers [23]. Furthermore, by moving away from using RA and administrative data points for predictive tasks, we can avoid abstracting ourselves to data points that distill and flatten social and economic problems in the public sector [6, 23, 101].

However, here we add a note of caution. Seidelin et al. [98] find that public sector algorithms are not always used the way they were initially intended, as is the case for RAs used in child welfare [91, 92]. Further, when technical interventions are repurposed from one context to the next they may cause unintended consequences [91, 93, 99] and automation bias [44, 95]. Just as AAPI and NCFAS scores were reappropriated to predict risk even though they were initially developed as support tools to help caseworkers identify appropriate services and goals for families [11, 64], it is always possible that new interventions may diverge from their original intentions. Recently, the private sector has introduced products that claim to harness NLP to assist child welfare caseworkers in case management [7, 8]. We argue that it is possible that such tools can diverge from their original intentions and be reappropriated amidst political and policy-related factors as has been done in the past [83, 91, 98]. In light of these concerns, we suggest researchers adopt multiple lenses involving both a quantitative and qualitative

lens to investigate stakeholder perceptions and the downstream effects of technical interventions.

8 CONCLUSION

In this study, we quantitatively deconstructed RA scores and child welfare narratives to assume the different roles of computing outlined by Abebe et al. [1]. We critique and add support to recent research that questions the construct validity of existing child welfare algorithms built on RA data by using casenotes and RA scores for cases in a US child welfare agency. In our study, we applied topic modeling to casenotes and trained random forest and support vector machines on RA scores and casenotes. We examined the predictive validity of these data sources on CW discharge outcomes. Our results show RAs could not inform discharge outcomes. We also show narrative casenotes contain caseworker perspectives on child welfare cases that are not uniform and may be biased, making casenotes unsuitable for case outcome prediction. We offer supporting evidence to move away from predictive tasks using RA data and instead find support for using contextual data such as casenotes to study sociotechnical systems in child welfare and, more broadly, in the public sector.

ACKNOWLEDGMENTS

This research was supported by the NSERC Discovery Early Career Researcher Grant RGPIN-2022-04570 and the Connaught New Researcher award. Any opinions, findings, conclusions, and recommendations expressed in this material are those of the authors. We are grateful for the anonymous reviewers whose suggestions and comments helped improve the quality of this manuscript.

REFERENCES

- [1] Rediet Abebe, Solon Barocas, Jon Kleinberg, Karen Levy, Manish Raghavan, and David G Robinson. 2020. Roles for computing in social change. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 252–260.
- [2] Doris Allhutter, Florian Cech, Fabian Fischer, Gabriel Grill, and Astrid Mager. 2020. Algorithmic profiling of job seekers in Austria: how austerity politics are made effective. *Front. Big Data* 3: 5. doi: 10.3389/fdata (2020).
- [3] Asbjørn Ammitzbøll Flügge, Thomas Hildebrandt, and Naja Holten Møller. 2021. Street-Level Algorithms and AI in Bureaucratic Decision-Making: A Caseworker Perspective. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW1 (2021), 1–23.
- [4] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. Machine Bias. Retrieved August 30, 2022 from <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- [5] Maria Antoniak, David Mimno, and Karen Levy. 2019. Narrative paths and negotiation of power in birth stories. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–27.
- [6] Cecilia Aragon, Shion Guha, Marina Kogan, Michael Muller, and Gina Neff. 2022. *Human-Centered Data Science: An Introduction*. MIT Press.
- [7] Augintel. 2020. About us: SOCIAL IMPACT MEETS A.I. Retrieved August 30, 2022 from <https://www.augintel.us/about>
- [8] Augintel. 2022. Allegheny County DHS Case Study: Unlocking the Data in Case Notes with Natural Language Processing. Retrieved September 4, 2022 from <https://www.youtube.com/watch?v=575pSgam6YY>
- [9] Solon Barocas and Andrew D Selbst. 2016. Big data's disparate impact. *Calif. L. Rev.* 104 (2016), 671.
- [10] Eric PS Baumer, David Mimno, Shion Guha, Emily Quan, and Geri K Gay. 2017. Comparing grounded theory and topic modeling: Extreme divergence or unlikely convergence? *Journal of the Association for Information Science and Technology* 68, 6 (2017), 1397–1410.
- [11] Stephen J. Bavolek and Richard G. Keene. 2010. AAPI Online Development Handbook The Adult-Adolescent Parenting Inventory (AAPI-2). *Family Development Resources, Inc* (2010). <https://www.nurturingparenting.com/images/aapionlinehandbook12-5-12.pdf>
- [12] S Bekaert, E Paavilainen, H Scheke, A Baldacchino, E Jouet, L Zablocka-Zytka, B Bachi, F Bartoli, G Carra, RM Cioni, et al. 2021. Family members' perspectives of

- child protection services, a metanalysis of the literature. *Children and Youth Services Review* (2021), 106094.
- [13] David M. Blei. 2012. Probabilistic Topic Models. *Commun. ACM* 55, 4 (apr 2012), 77–84. <https://doi.org/10.1145/2133806.2133826>
- [14] David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* 3 (March 2003), 993–1022.
- [15] Emily Adlin Bosk. 2018. What counts? Quantification, worker judgment, and divergence in child welfare decision making. *Human Service Organizations: Management, Leadership & Governance* 42, 2 (2018), 205–224.
- [16] Virginia Braun and Victoria Clarke. 2006. Using thematic analysis in psychology. *Qualitative Research in Psychology* 3, 2 (2006), 77–101. <https://doi.org/10.1191/1478088706qp063oa>
- [17] Leo Breiman. 2001. Random Forests. *Machine Learning* 45, 1 (01 Oct 2001), 5–32. <https://doi.org/10.1023/A:1010933404324>
- [18] Anna Brown, Alexandra Chouldechova, Emily Putnam-Hornstein, Andrew Tobin, and Rhema Vaithianathan. 2019. Toward Algorithmic Accountability in Public Services: A Qualitative Study of Affected Community Perspectives on Algorithmic Decision-making in Child Welfare Services. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. ACM, 41.
- [19] Jenna Burrell and Marion Fourcade. 2021. The Society of Algorithms. *Annual Review of Sociology* 47, 1 (2021), 213–237. <https://doi.org/10.1146/annurev-soc-090820-020800>
- [20] Peter D. Caie, Neofytos Dimitriou, and Ognjen Arandjelović. 2021. Chapter 8 - Precision medicine in digital pathology via image analysis and machine learning. In *Artificial Intelligence and Deep Learning in Pathology*, Stanley Cohen (Ed.). Elsevier, 149–173. <https://doi.org/10.1016/B978-0-323-67538-3.00008-7>
- [21] Michael J Camasso and Radha Jagannathan. 2013. Decision making in child protective services: A risky business? *Risk analysis* 33, 9 (2013), 1636–1649.
- [22] Capacity Building Center for States. 2018. Child Protective Services: A Guide for Caseworkers. <https://www.childwelfare.gov/pubPDFs/cps2018.pdf>
- [23] Stevie Chancellor, Eric P. S. Baumer, and Munmun De Choudhury. 2019. Who is the “Human” in Human-Centered Machine Learning: The Case of Predicting Mental Health from Social Media. *Proc. ACM Hum.-Comput. Interact.* 3, CSCW, Article 147 (nov 2019), 32 pages. <https://doi.org/10.1145/3359249>
- [24] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer. 2002. SMOTE: Synthetic Minority Over-sampling Technique. *Journal of Artificial Intelligence Research* 16 (jun 2002), 321–357. <https://doi.org/10.1613/jair.953>
- [25] Hao-Fei Cheng, Logan Stapleton, Anna Kawakami, Venkatesh Sivaraman, Yanghui Cheng, Diana Qing, Adam Perer, Kenneth Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. How Child Welfare Workers Reduce Racial Disparities in Algorithmic Decisions. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 162, 22 pages. <https://doi.org/10.1145/3491102.3501831>
- [26] Child Welfare Information Gateway. [n.d.]. Family Preservation Services. Retrieved September 2, 2022 from <https://www.childwelfare.gov/topics/supporting/preservation/>
- [27] Child Welfare Information Gateway. 2019. What Is Child Abuse and Neglect? Recognizing the Signs and Symptoms. Retrieved August 16, 2022 from <https://www.childwelfare.gov/pubPDFs/whatican.pdf>
- [28] Child Welfare Information Gateway. 2021. Concurrent Planning for Timely Permanency for Children. Retrieved September 2, 2022 from <https://www.childwelfare.gov/pubPDFs/concurrent.pdf>
- [29] Child Welfare Information Gateway. 2022. The use of safety and risk assessments in child protection cases. Retrieved August 16, 2022 from <https://www.childwelfare.gov/topics/systemwide/laws-policies/statutes/safety-risk/>
- [30] Alexandra Chouldechova, Diana Benavides-Prado, Oleksandr Fialko, and Rhema Vaithianathan. 2018. A case study of algorithm-assisted decision making in child maltreatment hotline screening decisions. In *Conference on Fairness, Accountability and Transparency*. 134–148.
- [31] Angèle Christin. 2017. Algorithms in practice: Comparing web journalism and criminal justice. *Big Data & Society* 4, 2 (2017), 2053951717718855. <https://doi.org/10.1177/2053951717718855>
- [32] Kevin Clancy, Joseph Chudzick, Aleksandra J. Snowden, and Shion Guha. 2022. Reconciling data-driven crime analysis with human-centered algorithms. *Cities* 124 (2022), 103604. <https://doi.org/10.1016/j.cities.2022.103604>
- [33] Ontario Human Rights Commission. 2017. Under suspicion: Concerns about child welfare. Retrieved September 2, 2022 from <https://www.ohrc.on.ca/en/under-suspicion-concerns-about-child-welfare>
- [34] Roxanne Connelly, Christopher J Playford, Vernon Gayle, and Chris Dibben. 2016. The role of administrative data in the big data revolution in social science research. *Social science research* 59 (2016), 1–12.
- [35] Nicola A. Conners, Leanne Whiteside-Mansell, David Deere, Toni Ledet, and Mark C. Edwards. 2006. Measuring the potential for child maltreatment: The reliability and validity of the Adult Adolescent Parenting Inventory–2. *Child Abuse & Neglect* 30, 1 (2006), 39–53. <https://doi.org/10.1016/j.chiabu.2005.08.011>
- [36] Victoria A Copeland. 2021. “It’s the Only System We’ve Got”: Exploring Emergency Response Decision-Making in Child Welfare. *Columbia Journal of Race and Law* 11, 3 (2021), 43–74.
- [37] Corinna Cortes and Vladimir Vapnik. 1995. Support-vector networks. *Machine Learning* 20, 3 (01 Sep 1995), 273–297. <https://doi.org/10.1007/BF00994018>
- [38] Amanda Coston, Anna Kawakami, Haiyi Zhu, Ken Holstein, and Hoda Heidari. 2023. A Validity Perspective on Evaluating the Justified Use of Data-driven Decision-making Algorithms. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*. IEEE, 690–704.
- [39] Dipto Das, Shion Guha, and Bryan Semaan. 2023. Toward Cultural Bias Evaluation Datasets: The Case of Bengali Gender, Religious, and National Identity. In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, Sunipa Dev, Vinodkumar Prabhakaran, David Adelani, Dirk Hovy, and Luciana Benotti (Eds.). Association for Computational Linguistics, Dubrovnik, Croatia, 68–83. <https://doi.org/10.18653/v1/2023.c3nlp-1.8>
- [40] Alan J. Dettlaff and Reiko Boyd. 2020. Racial Disproportionality and Disparities in the Child Welfare System: Why Do They Exist, and What Can Be Done to Address Them? *The ANNALS of the American Academy of Political and Social Science* 692, 1 (2020), 253–274. <https://doi.org/10.1177/0002716220980329>
- [41] Alan J Dettlaff, Stephanie L Rivaux, Donald J Baumann, John D Fluke, Joan R Rycraft, and Joyce James. 2011. Disentangling substantiation: The influence of race, income, and risk on the substantiation decision in child welfare. *Children and Youth Services Review* 33, 9 (2011), 1630–1637.
- [42] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science Advances* 4, 1 (2018), eao5580. <https://doi.org/10.1126/sciadv.aao5580>
- [43] Edx. 2015. Adult-Adolescent Parenting Inventory-2 (AAPI-2). <https://edge.edx.org/assets/courseware/v1/cf4d30fc491e49772c24e144cf3c0c3/asset-v1:GeorgetownX+CCHD+2016+type@asset-block/AAPI.pdf>
- [44] Virginia Eubanks. 2018. *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin’s Press.
- [45] Family Development Resources. [n. d.]. Inventory Scoring System for Assessing Parenting Practices. Retrieved May 30, 2023 from <https://assessingparenting.com>
- [46] Anjalie Field, Amanda Coston, Nupoor Gandhi, Alexandra Chouldechova, Emily Putnam-Hornstein, David Steier, and Yulia Tsvetkov. 2023. Examining Risks of Racial Biases in NLP Tools for Child Protective Services. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1479–1492. <https://doi.org/10.1145/3593013.3594094>
- [47] Annie Franco, Neil Malhotra, and Gabor Simonovits. 2014. Publication bias in the social sciences: Unlocking the file drawer. *Science* 345, 6203 (2014), 1502–1505. <https://doi.org/10.1126/science.1255484> arXiv:<https://www.science.org/doi/pdf/10.1126/science.1255484>
- [48] Mikel Galar, Alberto Fernandez, Edurne Barrenechea, Humberto Bustince, and Francisco Herrera. 2012. A Review on Ensembles for the Class Imbalance Problem: Bagging-, Boosting-, and Hybrid-Based Approaches. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 42, 4 (2012), 463–484. <https://doi.org/10.1109/TSMCC.2011.2161285>
- [49] Jennifer M Geiger and Lisa Schelbe. 2021. Foster Care Placement. In *The Handbook on Child Welfare Practice*. Springer, 219–248.
- [50] Marissa Gerchick, Tobi Jegede, Tarak Shah, Ana Gutierrez, Sophie Beiers, Noam Shemtov, Kath Xu, Anjana Samant, and Aaron Horowitz. 2023. The Devil is in the Details: Interrogating Values Embedded in the Allegheny Family Screening Tool. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1292–1310. <https://doi.org/10.1145/3593013.3594081>
- [51] Philip Gillingham. 2018. Decision-making about the adoption of information technology in social welfare agencies: Some key considerations. *European Journal of Social Work* 21, 4 (2018), 521–529.
- [52] Jeremy D Goldhaber-Fiebert and Lea Prince. 2019. *Impact Evaluation of a Predictive Risk Modeling Tool for Allegheny County’s Child Welfare Office*. Technical Report.
- [53] Travis Greene, Galit Shmueli, Jan Fell, Ching-Fu Lin, Mark L Shope, and Han-Wei Liu. 2020. The Hidden Inconsistencies Introduced by Predictive Algorithms in Judicial Decision Making. *arXiv preprint arXiv:2012.00289* (2020).
- [54] Sally Ho and Garance Burke. 2022. An algorithm that screens for child neglect raises concerns. <https://apnews.com/article/child-welfare-algorithm-investigation-9497ee937e0053ad4144a86c68241ef1>
- [55] Naja Holten Nielsen, Trine Rask Nielsen, and Christopher Le Dantec. 2021. Work of the Unemployed: An Inquiry into Individuals’ Experience of Data Usage in Public Services and Possibilities for Their Agency. In *Designing Interactive Systems Conference 2021* (Virtual Event, USA) (DIS '21). Association for Computing Machinery, New York, NY, USA, 438–448. <https://doi.org/10.1145/3461778.3462003>
- [56] Naja Holten Nielsen, Irina Shklovski, and Thomas T Hildebrandt. 2020. Shifting concepts of value: Designing algorithmic decision-support systems for public services. In *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*. 1–12.

- [57] Hornby Zeller Associates INC. 2018. *Allegheny County Predictive Risk Modeling Tool Implementation: Process Evaluation*. Technical Report.
- [58] Abigail Z. Jacobs, Su Lin Blodgett, Solon Barocas, Hal Daumé, and Hanna Wallach. 2020. The Meaning and Measurement of Bias: Lessons from Natural Language Processing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (FAT* '20). Association for Computing Machinery, New York, NY, USA, 706. <https://doi.org/10.1145/3351095.3375671>
- [59] Nathalie Japkowicz and Shaju Stephen. 2002. The class imbalance problem: A systematic study. *Intelligent Data Analysis* (2002), 429–449.
- [60] Anna Kawakami, Venkatesh Sivaraman, Hao-Fei Cheng, Logan Stapleton, Yanghui Cheng, Diana Qing, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. Improving Human-AI Partnerships in Child Welfare: Understanding Worker Practices, Challenges, and Desires for Algorithmic Decision Support. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems* (New Orleans, LA, USA) (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 52, 18 pages. <https://doi.org/10.1145/3491102.3517439>
- [61] Anna Kawakami, Venkatesh Sivaraman, Logan Stapleton, Hao-Fei Cheng, Adam Perer, Zhiwei Steven Wu, Haiyi Zhu, and Kenneth Holstein. 2022. “Why Do I Care What’s Similar?” Probing Challenges in AI-Assisted Child Welfare Decision-Making through Worker-AI Interface Design Concepts. In *Designing Interactive Systems Conference* (Virtual Event, Australia) (DIS '22). Association for Computing Machinery, New York, NY, USA, 454–470. <https://doi.org/10.1145/3532106.3533556>
- [62] Emily Keddell. 2015. The ethics of predictive risk modelling in the Aotearoa/New Zealand child welfare context: Child abuse prevention or neo-liberal tool? *Critical Social Policy* 35, 1 (2015), 69–88. <https://doi.org/10.1177/0261018314543224>
- [63] Raymond S. Kirk. 2015. Psychometric Properties of the Trauma and Post-Trauma Well-Being Assessment Domains of the North Carolina Family Assessment Scale for General and Reunification Services (NCFAS G+R). *Journal of Public Child Welfare* 9, 5 (2015), 444–462. <https://doi.org/10.1080/15548732.2015.1090364>
- [64] Raymond S. Kirk and Kellie Reed-Ashcraft. 2004. NCFAS North Carolina Family Assessment Scale Research Report. *National Family Preservation Network* (2004), 1–22. https://www.nfpn.org/media/8d86bbe82c2a88d/ncfas_research_report.pdf
- [65] Brianne H Kothari, Kelly D Chandler, Andrew Waugh, Kara K McElvaine, Jamie Jaramillo, and Shannon Lipscomb. 2021. Retention of child welfare caseworkers: The role of case severity and workplace resources. *Children and Youth Services Review* 126 (2021), 106039.
- [66] Kweku Kwegyir-Aggrey, Marissa Gerchick, Malika Mohan, Aaron Horowitz, and Suresh Venkatasubramanian. 2023. The Misuse of AUC: What High Impact Risk Assessment Gets Wrong. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) (FAccT '23). Association for Computing Machinery, New York, NY, USA, 1570–1583. <https://doi.org/10.1145/3593013.3594100>
- [67] Bethany R. Lee and Michael A. Lindsey. 2010. North Carolina Family Assessment Scale: Measurement Properties for Youth Mental Health Services. *Research on Social Work Practice* 20, 2 (2010), 202–211. <https://doi.org/10.1177/1049731509334180>
- [68] Yaguo Lei. 2017. 3 - Individual intelligent method-based fault diagnosis. In *Intelligent Fault Diagnosis and Remaining Useful Life Prediction of Rotating Machinery*, Yaguo Lei (Ed.). Butterworth-Heinemann, 67–174. <https://doi.org/10.1016/B978-0-12-811534-3.00003-2>
- [69] Guillaume Lemaître, Fernando Nogueira, and Christos K. Aridas. 2017. Imbalanced-learn: A Python Toolbox to Tackle the Curse of Imbalanced Datasets in Machine Learning. *Journal of Machine Learning Research* 18, 17 (2017), 1–5. <http://jmlr.org/papers/v18/16-365.html>
- [70] Jorge M. Lobo, Alberto Jiménez-Valverde, and Raimundo Real. 2008. AUC: A misleading measure of the performance of predictive distribution models. *Global Ecology and Biogeography* 17, 2 (2008), 145–151. <https://doi.org/10.1111/j.1466-8238.2007.00358.x>
- [71] Paul Marks. 2022. Algorithmic Hiring Needs a Human Face. *Commun. ACM* 65, 3 (feb 2022), 17–19. <https://doi.org/10.1145/3510552>
- [72] Kelly McConvey, Shion Guha, and Anastasia Kuzminykh. 2023. A Human-Centered Review of Algorithms in Decision-Making in Higher Education. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (Hamburg, Germany) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 223, 15 pages. <https://doi.org/10.1145/3544548.3580658>
- [73] Nicole Mickelson, Traci LaLiberte, and Kristine Piescher. 2017. Assessing Risk: A Comparison of Tools for Child Welfare Practice with Indigenous Families. Retrieved June 15, 2023 from <https://casew.umn.edu/portfolio-items/assessing-risk-a-comparison-of-tools-for-child-welfare-practice-with-indigenous-families/>
- [74] Erina Seh Young Moon and Shion Guha. 2024. A Human-Centered Review of Algorithms in Homelessness Research. (2024). <https://doi.org/10.1145/3613904.3642392> arXiv:arXiv:2401.13247
- [75] Ramaravind Kommiya Mothilal, Shion Guha, and Syed Ishtiaque Ahmed. 2024. Towards a Non-Ideal Methodological Framework for Responsible ML. arXiv:2401.11131 [cs.HC]
- [76] Michael Muller, Shion Guha, Eric PS Baumer, David Mimno, and N Sadat Shami. 2016. Machine learning and grounded theory method: Convergence, divergence, and combination. In *Proceedings of the 19th International Conference on Supporting Group Work*. ACM, 3–8.
- [77] National Family Preservation Network. 2009. NCFAS North Carolina Family Assessment Scale Scale & Definitions (v. 2.0). (2009). Retrieved September 2, 2022 from https://cscbroward.org/sites/default/files/2018-06/NCFAS%20Scale_Defs.pdf
- [78] Dong Nguyen, Maria Liakata, Simon DeDeo, Jacob Eisenstein, David Mimno, Rebekah Tromble, and Jane Winters. 2020. How We Do Things With Words: Analyzing Text as Social and Cultural Data. *Frontiers in Artificial Intelligence* 3 (2020), 62.
- [79] Lan Huong Nguyen and Susan Holmes. 2019. Ten quick tips for effective dimensionality reduction. *JPLoS Comput Biol*. 15, 6, Article 6:e1006907 (june 2019). <https://doi.org/10.1371/journal.pcbi.1006907>
- [80] William S. Noble. 2006. What is a support vector machine? *Nature Biotechnology* 24, 12 (01 Dec 2006), 1565–1567. <https://doi.org/10.1038/nbt1206-1565>
- [81] Yotam Ophir, Dror Walter, and Eleanor R Marchant. 2020. A Collaborative Way of Knowing: Bridging Computational Communication Research and Grounded Theory Ethnography. *Journal of Communication* 70, 3 (2020), 447–472.
- [82] Inioluwa Deborah Raji, I. Elizabeth Kumar, Aaron Horowitz, and Andrew Selbst. 2022. The Fallacy of AI Functionality. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (FAccT '22). Association for Computing Machinery, New York, NY, USA, 959–972. <https://doi.org/10.1145/3531146.3533158>
- [83] Joanna Redden, Lina Dencik, and Harry Warne. 2020. Datafied child welfare services: Unpacking politics, economics and power. *Policy Studies* 41, 5 (2020), 507–526.
- [84] Cheryl Regehr, Marion Bogo, Aron Shlonsky, and Vicki LeBlanc. 2010. Confidence and Professional Judgment in Assessing Children’s Risk of Abuse. *Research on Social Work Practice* 20, 6 (2010), 621–628. <https://doi.org/10.1177/1049731510368050>
- [85] Dorothy E. Roberts. 2002. RACIAL HARM: Dorothy Roberts explains how racism works in the child welfare system. *Colorlines* 5, 3 (Fall 2002), 19.
- [86] Dorothy E. Roberts. 2022. *Torn Apart How the Child Welfare System Destroys Black Families—and How Abolition Can Build a Safer World*. Basic Books, New York, NY, USA.
- [87] Samantha Robertson, Tonya Nguyen, and Niloufar Salehi. 2021. Modeling Assumptions Clash with the Real World: Transparency, Equity, and Community Challenges for Student Assignment Algorithms. *arXiv preprint arXiv:2101.10367* (2021).
- [88] Samantha Robertson and Niloufar Salehi. 2020. What If I Don’t Like Any Of The Choices? The Limits of Preference Elicitation for Participatory Algorithm Design. *arXiv preprint arXiv:2007.06718* (2020).
- [89] Anjana Samant, Aaron Horowitz, Sophie Beiers, and Kath Xu. 2021. Family Surveillance by Algorithm: The Rapidly Spreading Tools Few Have Heard Of. Retrieved May 2, 2023 from <https://www.aclu.org/news/womens-rights/family-surveillance-by-algorithm-the-rapidly-spreading-tools-few-have-heard-of>
- [90] Devansh Saxena, Karla Badillo-Urquiola, Pamela Wisniewski, and Shion Guha. 2020. Child welfare system: Interaction of policy, practice and algorithms. In *Companion Proceedings of the 2020 ACM International Conference on Supporting Group Work*. 119–122.
- [91] Devansh Saxena, Karla Badillo-Urquiola, Pamela Wisniewski, and Shion Guha. 2021. A Framework of High-Stakes Algorithmic Decision-Making for the Public Sector Developed through a Case Study of Child-Welfare. *Proceedings of the ACM on Human-Computer Interaction* 5, CSCW2 (2021).
- [92] Devansh Saxena, Karla Badillo-Urquiola, Pamela J Wisniewski, and Shion Guha. 2020. A Human-Centered Review of Algorithms used within the US Child Welfare System. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–15.
- [93] Devansh Saxena and Shion Guha. 2023. Algorithmic Harms in Child Welfare: Uncertainties in Practice, Organization, and Street-level Decision-Making. *ACM J. Responsib. Comput.* (sep 2023). <https://doi.org/10.1145/3616473>
- [94] Devansh Saxena, Erina Seh-Young Moon, Aryan Chaurasia, Yixin Guan, and Shion Guha. 2023. Rethinking “Risk” in Algorithmic Systems Through A Computational Narrative Analysis of Casenotes in Child-Welfare. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–19.
- [95] Devansh Saxena, Seh Young Moon, Dahlia Shehata, and Shion Guha. 2022. Unpacking Invisible Work Practices, Constraints, and Latent Power Relationships in Child Welfare through Casenote Analysis (CHI '22). Association for Computing Machinery, New York, NY, USA, Article 120, 22 pages. <https://doi.org/10.1145/3491102.3517742>
- [96] Devansh Saxena, Charles Repaci, Melanie D Sage, and Shion Guha. 2022. How to Train a (Bad) Algorithmic Caseworker: A Quantitative Deconstruction of Risk Assessments in Child Welfare. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts*. 1–7.
- [97] Craig S Schwalbe. 2008. Strengthening the integration of actuarial risk assessment with clinical judgment in an evidence based practice framework. *Children*

- and Youth Services Review* 30, 12 (2008), 1458–1464.
- [98] Cathrine Seidelin, Therese Moreau, Irina Shklovski, and Naja Holten Møller. 2022. Auditing Risk Prediction of Long-Term Unemployment. *Proc. ACM Hum.-Comput. Interact.* 6, GROUP, Article 8 (jan 2022), 12 pages. <https://doi.org/10.1145/3492827>
- [99] Andrew D. Selbst, Danah Boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2019. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (Atlanta, GA, USA) (*FAT* '19*). Association for Computing Machinery, New York, NY, USA, 59–68. <https://doi.org/10.1145/3287560.3287598>
- [100] Aron Shlonsky and Dennis Wagner. 2005. The next step: Integrating actuarial risk assessment and clinical judgment into an evidence-based practice framework in CPS case management. *Children and youth services review* 27, 4 (2005), 409–427.
- [101] Ben Shneiderman. 2022. *Human-Centered AI*. Oxford University Press. <https://doi.org/10.1093/oso/9780192845290.001.0001>
- [102] Nancy L. Sidell. 2015. *Social Work Documentation*. NASW Press, Washington, DC, USA.
- [103] Gabriel Tobin Smith, Valerie B Shapiro, Rachel Wagner Sperry, and Paul A LeBuffe. 2014. A strengths-based approach to supervised visitation in child welfare. *Child Care in Practice* 20, 1 (2014), 98–119.
- [104] Logan Stapleton, Min Hun Lee, Diana Qing, Marya Wright, Alexandra Chouldechova, Ken Holstein, Zhiwei Steven Wu, and Haiyi Zhu. 2022. Imagining new futures beyond predictive systems in child welfare: A qualitative study with impacted stakeholders. In *FAccT '22: 2022 ACM Conference on Fairness, Accountability, and Transparency, Seoul, Republic of Korea, June 21 - 24, 2022*. ACM, 1162–1177. <https://doi.org/10.1145/3531146.3533177>
- [105] Zach Strassburger. 2016. Medical Decision Making for Youth in the Foster Care System. *J. Marshall L. Rev.* 49, 4, Article 5 (2016), 1103–1154 pages. <https://repository.law.uic.edu/lawreview/vol49/iss4/5/>
- [106] Dana Goldstein. New York Times. April 2019. San Francisco Had an Ambitious Plan to Tackle School Segregation. It Made It Worse. <https://www.nytimes.com/2019/04/25/us/san-francisco-school-segregation.html>
- [107] Rhema Vaithianathan, Diana Benavides-Prado, Erin Dalton, Alex Chouldechova, and Emily Putnam-Hornstein. 2021. Using a Machine Learning Tool to Support High-Stakes Decisions in Child Protection. *AI Magazine* 42, 1 (2021), 53–60. <https://ojs.aaai.org/index.php/aimagazine/article/view/7482>
- [108] Rhema Vaithianathan, Emily Putnam-Hornstein, Alexandra Chouldechova, Diana Benavides-Prado, and Rachel Berger. 2020. Hospital Injury Encounters of Children Identified by a Predictive Risk Model for Screening Child Maltreatment Referrals: Evidence From the Allegheny Family Screening Tool. *JAMA pediatrics* 174, 11 (Nov 2020), e202770. <https://doi.org/10.1001/jamapediatrics.2020.2770>
- [109] Rhema Vaithianathan, Emily Putnam-Hornstein, Nan Jiang, Parma Nand, and Tim Maloney. 2017. Developing predictive models to support child maltreatment hotline screening decisions: Allegheny County methodology and implementation. *Center for Social data Analytics* (2017).
- [110] Michael Veale and Irina Brass. 2019. Administration by Algorithm? Public Management Meets Public Sector Machine Learning. 3375391 (2019). <https://papers.ssrn.com/abstract=3375391>
- [111] Michael Veale, Max Van Kleek, and Reuben Binns. 2018. Fairness and accountability design needs for algorithmic support in high-stakes public sector decision-making. In *Proceedings of the 2018 chi conference on human factors in computing systems*. 1–14.
- [112] Maranke Wieringa. 2020. What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (Barcelona, Spain) (*FAT* '20*). Association for Computing Machinery, New York, NY, USA, 1–18. <https://doi.org/10.1145/3351095.3372833>
- [113] Ben Williamson. 2016. Digital education governance: Data visualization, predictive analytics, and 'real-time' policy instruments. *Journal of Education Policy* 31, 2 (2016), 123–141.
- [114] Max L. Wilson, Wendy Mackay, Ed Chi, Michael Bernstein, Dan Russell, and Harold Thimbleby. 2011. RepliCHI - CHI Should Be Replicating and Validating Results More: Discuss. In *CHI '11 Extended Abstracts on Human Factors in Computing Systems* (Vancouver, BC, Canada) (*CHI EA '11*). Association for Computing Machinery, New York, NY, USA, 463–466. <https://doi.org/10.1145/1979742.1979491>